

Using CORBA and Speech Recognition for Structured Medical Reporting

David Rosenthal and Jennifer Mayo
Kurzweil Applied Intelligence, Waltham Massachusetts

Background. Clinical reporting represents a major cost to the healthcare industry, in terms of both dollars and time. The transcription services used by many physicians are expensive and error-prone. The output of the usual reporting process is either ASCII, or worse, paper documents, which, while they fulfill the traditional purpose of documenting patient encounters, are difficult or impossible to integrate with other records of the (increasingly electronic) longitudinal patient record.

Structured reporting systems can address some of these issues – in particular, such systems can produce electronic documents whose contents may be automatically accessed by other local or remote information systems. Structured reporting is difficult to learn, and this learning curve tends to require even more of the physician's time than traditional reporting by transcription.

Kurzweil AI has received a grant from the National Institute of Standards and Technology's Advanced Technology Program to develop a voice-enabled, structured reporting system which produces output in industry-standard formats, and is therefore accessible by heterogeneous systems and networks. Our aim, in effect, is to solve both of the above-mentioned problems simultaneously. In ease of use, our system is comparable to traditional reporting methods, yet the output is a structured electronic document.

System. To use the system, physicians normally dictate reports from informal notes taken during the encounter with the patient. The dictation session is mediated by a knowledge base of clinical information which guides the construction of a structured report from the physician's utterances. The output is represented in two industry standard formats: a text format called SGML (Standard Generalized Markup Language), and as a distributed *object* in a binary format called CORBA (Common Object Request Broker Architecture). CORBA is an industry standard for *distributed object computing*. Distributed objects are accessible to local or remote clients which access them through *methods*. In our case, the report is represented as an object, and typical methods are those which

get or set sections of the report – so for example, one method might update the patient's current medications; another might retrieve the patient's vital signs.

Our currently implemented system uses CORBA to communicate with a sample database of patient records maintained at CareFlowNet, a healthcare and Internet products firm in West Virginia.

Standards Development. Although CORBA is a computing industry standard, the healthcare industry needs to develop a particular protocol for communicating with the clinical report object. This protocol is defined in a document called an IDL (Interface Definition Language). To facilitate communication among diverse reporting systems that produce CORBA objects, the IDL should itself be standardized. Kurzweil AI is an active participant in CORBAmed, an organization whose purpose is to standardize healthcare-related CORBA objects and IDLs.

Conclusions. Physicians are currently in a difficult position: On one hand, they are under pressure to cut costs; on the other hand, their documentation requirements are increasing. Traditional paper records of patient encounters are now considered inadequate in many instances; there is an increasing demand that records be in an electronic format and available, in a standard format, across a network connection. We believe that our system can address all of these issues. A voice-enabled, structured reporting system gives the physician a means to produce reports quickly, efficiently, and much more cheaply than those currently in widespread use. At the same time, the ability to produce structured output in a standard format which is available over the Internet and easily accessed for reporting and analysis makes the information generated by physicians much more useful and valuable.